5

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/624,242	07/21/2003	Sashikanth Chandrasekaran	ORCL.P102	7566

42425 7590 12/20/2006

HICKMAN PALERMO TRUONG & BECKER/ORACLE
2055 GATEWAY PLACE
SUITE 550
SAN JOSE, CA 95110-1089

EXAMINER	
CHU, GABRIEL L	
ART UNIT	**PAPER NUMBER**
2114	

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
3 MONTHS	12/20/2006	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary	Application No.	Applicant(s)	
	10/624,242	CHANDRASEKARAN, SASHIKANTH	
	Examiner	Art Unit	
	Gabriel L. Chu	2114	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *21 July 2003*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-24* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-24* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *21 July 2003* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *0222 0215*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

1. Claim 7 objected to because of the following informalities:

Referring to claim 7, "until another a" is understood to refer to "until a".

Appropriate correction is required.

Claim Rejections - 35 USC § 101

2. 35 U.S.C. 101 reads as follows:

> Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

Claims 21, 22 rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter. Referring to claims 21, 22, the claims have a computer-readable medium "carrying" instructions. From paragraph 78 of the pre-grant publication, such mediums are disclosed to include non-statutory transmission media. To overcome this rejection, the claims must be amended to state that the mediums are "storing" the instructions.

Claim Rejections - 35 USC § 102

3. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

4. Claims 1-9 14-16, 18, 21-24 rejected under 35 U.S.C. 102(b) as being anticipated by US 5550973 to Forman et al.

5. Referring to claim 1, 21, 23, Forman discloses a method of operation within a data processing system that includes a plurality of processing nodes each having access to a set of shared resources (From line 9 of column 1, "The present invention relates to the operation of distributed processing computer systems. In particular it relates to those systems that have a plurality of processing nodes each one having access to a copy of at least one replicated data object and which require apparatus and methods for managing access to the replicated data object. Still more particularly, the present invention relates to the management of a write lock that grants permission to one of a number of distributed processes allowing that process to update a data item."), the method comprising:

detecting failure of a failed node of the plurality of processing nodes (From line 7 of column 4, "Another advantage is that failure of any particular node will not cause the entire network to fail as would be the case where a master processor existed." Wherein Forman discloses any node of the nodes can fail.);

receiving a request to access a first resource of the set of shared resources (Figure 3, element 150.);

and granting access to the first resource if the failed node was not responsible for controlling access to the first resource and did not have exclusive access to the first resource when the failure was detected (Wherein the node that fails is neither the master nor the write lock holder, this is determined by not detecting a failure of the master, and in either case (Figure 4, element 202, figure 3, element 176), if the failed node is also not the lock holder, access is granted (Figure 3, element 174). Further,

from line 37 of column 4, "The master process has the responsibility for managing a write lock that permits the holder to update its copy of the replicated data object. Any process desiring to update a data object first requests the write lock from the master process. If the master does not have the lock, it determines which process does and then sends a message to that process to forward the lock to the requesting process.").

6. Referring to claim 2, Forman discloses granting access to the first resource comprises: determining whether the failed node was responsible for controlling access to the first resource; and if the failed node was not responsible for controlling access to the first resource, determining whether, at the time the failure was detected, the failed node had exclusive access to the first resource (Figure 4, element 202, figure 3, element 176), if the failed node is also not the lock holder, access is granted (Figure 3, element 174). Further, from line 37 of column 4, "The master process has the responsibility for managing a write lock that permits the holder to update its copy of the replicated data object. Any process desiring to update a data object first requests the write lock from the master process. If the master does not have the lock, it determines which process does and then sends a message to that process to forward the lock to the requesting process.").

7. Referring to claim 3, Forman discloses determining whether the failed node was responsible for controlling access to the first resource comprises inspecting a data structure that indicates, for each shared resource within the set of shared resources, which of the plurality of processing nodes is responsible for controlling access to the shared resource (From line 23 of column 4, "If the attempt to acquire exclusive write

lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170." From line 27 of column 5, "The shared control file according the preferred embodiment of the present invention, contains the process identity and communications address of the write lock holder.").

8. Referring to claim 4, Forman discloses determining whether the failed node was responsible for controlling access to the first resource comprises identifying a data element within the data structure that includes a first component that identifies the first resource and a second component that identifies a processing node responsible for controlling access to the first resource (Figure 3, element 152. From line 23 of column 4, "If the attempt to acquire exclusive write lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170." From line 27 of column 5, "The shared control file according the preferred embodiment of the present invention, contains the process identity and communications address of the write lock holder.").

9. Referring to claim 5, Forman discloses preventing access to the first resource if the failed node was responsible for controlling access to the first resource or if the failed node had exclusive access to the first resource when the failure was detected (Figure 4, wherein the master changing process prevents access for at least the period wherein the master is changing, as no access may be granted during that period.).

10. Referring to claim 6, 16, Forman discloses preventing access comprises denying access to the first resource (Figure 4, wherein the master changing process prevents

access for at least the period wherein the master is changing, as no access may be granted during that period.).

11. Referring to claim 7, Forman discloses preventing access comprises deferring access to the first resource at least until another a first surviving node of the plurality of processing nodes is assigned responsibility for the first resource and the first surviving node determines whether, at the time the failure was detected, any of the plurality of processing nodes other than the failed node had access to the first resource (Figure 4, all, but also specifically element 202.).

12. Referring to claim 8, Forman discloses granting access to the first resource if the failed node was responsible for controlling access to the first resource and if a surviving node of the plurality of processing nodes had access to the first resource when the failure was detected (Figure 4, element 204. From line 23 of column 4, "If the attempt to acquire exclusive write lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170. If the requesting process is the master, it can directly access the resource, otherwise, it is a shadow process and must negotiate with the master for access 176.").

13. Referring to claim 9, 18, Forman discloses reassigning responsibility for controlling access to the first resource to a first surviving node of the plurality of processing nodes if the failed node was responsible for controlling access to the first resource (Figure 4).

14. Referring to claim 14, 22, 24, Forman discloses a method of operation within a data processing system that includes a plurality of processing nodes each having access to a set of shared resources (From line 9 of column 1, "The present invention relates to the operation of distributed processing computer systems. In particular it relates to those systems that have a plurality of processing nodes each one having access to a copy of at least one replicated data object and which require apparatus and methods for managing access to the replicated data object. Still more particularly, the present invention relates to the management of a write lock that grants permission to one of a number of distributed processes allowing that process to update a data item."), the method comprising:

detecting failure of a failed node of the plurality of processing nodes (From line 7 of column 4, "Another advantage is that failure of any particular node will not cause the entire network to fail as would be the case where a master processor existed." Wherein Forman discloses any node of the nodes can fail.);

receiving a request to access a first resource of the set of shared resources (Figure 3, element 150.);

and granting access to the first resource if (i) the failed node was responsible for controlling access to the first resource and (ii) a non-failed node of the plurality of processing nodes had access to the first resource when the failure was detected (Figure 4, element 204. From line 23 of column 4, "If the attempt to acquire exclusive write lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170. If the

requesting process is the master, it can directly access the resource, otherwise, it is a shadow process and must negotiate with the master for access 176.").

15. Referring to claim 15, Forman discloses preventing access to the first resource if the failed node was responsible for controlling access to the first resource and none of the non-failed nodes of the plurality of processing nodes had access to the first resource when the failure was detected (Figure 4, element 204, wherein the master changing process prevents access for at least the period wherein the master is changing, as no access may be granted during that period.).

Claim Rejections - 35 USC § 103

16. Claim 10, 11, 19, 20 rejected under 35 U.S.C. 103(a) as being unpatentable over US 5550973 to Forman et al., as applied to claims 9, 18 above.

17. Referring to claim 10, 19, Forman discloses generating a data structure that indicates whether a processing node of the plurality of processing nodes, other than the failed node, had access to the first resource when the failure was detected (Figure 3, element 154, 152. From line 23 of column 4, "If the attempt to acquire exclusive write lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170." From line 27 of column 5, "The shared control file according the preferred embodiment of the present invention, contains the process identity and communications address of the write lock holder.").

Although Forman does not explicitly disclose in the detailed embodiments that

the storage control file is within the first surviving node, having a file to the manipulating agent was known by Forman. From the background line 28 of column 1, "A replicated object is a logical unit of data existing in one of the computer systems but physically replicated to multiple distributed computer systems. Replicated copies are typically maintained in the memories of the distributed systems. Replicated data objects also speed the update process by allowing immediate local update of a data object. Replication introduces a control problem, however, because many copies of the data object exist. The distributed system must have some means for controlling data update to ensure that all copies of the data remain consistent." Forman, or a person having ordinary skill in the art, would have been motivated to have the SCF within the surviving node because "Replicated data objects also speed the update process by allowing immediate local update of a data object", and since the surviving/shadow node becomes the master, consistency is maintained.

18. Referring to claim 11, Forman discloses granting access to the first resource if (i) responsibility for controlling access to the first resource was reassigned after the failure was detected and (ii) the data structure indicates that a processing node, other than the failed node, had access to the first resource when the failure was detected (Figure 4, element 204. From line 23 of column 4, "If the attempt to acquire exclusive write lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170. If the requesting process is the master, it can directly access the resource, otherwise, it is a shadow process and must negotiate with the master for access 176.").

19. Referring to claim 20, Forman discloses granting access to the first resource comprises granting access to the first resource after reassigning responsibility for controlling access to the first resource to the first non-failed node and after generating the data structure that indicates whether a processing node other than the failed node had access to the first resource (Figure 4, element 204. From line 23 of column 4, "If the attempt to acquire exclusive write lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170. If the requesting process is the master, it can directly access the resource, otherwise, it is a shadow process and must negotiate with the master for access 176.").

20. Claim 12, 13 rejected under 35 U.S.C. 103(a) as being unpatentable over US 5550973 to Forman et al. as applied to claim 1 above, and further in view of US 6374362 to Ohtsu.

21. Referring to claim 12, Forman discloses identifying which node is responsible for controlling access and which node has exclusive access (Figure 4, element 204. From line 23 of column 4, "If the attempt to acquire exclusive write lock failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170. If the requesting process is the master, it can directly access the resource, otherwise, it is a shadow process and must negotiate with the master for access 176.").

Although Forman does not specifically disclose adding an identifier of the first resource to a validation data structure if the failed node was not responsible for

controlling access to the first resource but had exclusive access to the first resource when the failure was detected, indicating which shadow process had access at failure is known in the art. An example of this is shown by Ohtsu, from the abstract, "At occurrence of a failure on a "failure" node within the multiple nodes, the node failure recovery unit of a certain "normal" node clears the node flag while placing the process control block in a state to wait for a service request." A person of ordinary skill in the art at the time of the invention would have been motivated to use such an indicator, as shown by Ohtsu, for recovery and subsequent access arbitration. Further, as shown by Forman, there is a need to determine accessibility for access.

22. Referring to claim 13, Forman in view of Ohtsu discloses granting access to the first resource comprises granting access to the first resource if (i) the failed node was not responsible for controlling access to the first resource and (ii) the validation data structure does not include an identifier of the first resource (Figure 4, element 202, figure 3, element 176), if the failed node is also not the lock holder, access is granted (Figure 3, element 174). Further, from line 37 of column 4, "The master process has the responsibility for managing a write lock that permits the holder to update its copy of the replicated data object. Any process desiring to update a data object first requests the write lock from the master process. If the master does not have the lock, it determines which process does and then sends a message to that process to forward the lock to the requesting process.").

23. Claim 17 rejected under 35 U.S.C. 103(a) as being unpatentable over US 5550973 to Forman et al. as applied to claim 15 above, and further in view of US 5920872 to Grewell et al.

24. Referring to claim 17, Forman discloses deferring access (Figure 4, element 204, wherein the master changing process prevents access for at least the period wherein the master is changing, as no access may be granted during that period.) and that there may be one or more transactions recorded in a failed master that have failed to propagate (From line 55 of column 4, "This could occur where a shadow process has commenced update but has not completed the process of sending the update to the master for propagation or where the master failed before propagating the data.").

Although Forman does not specifically disclose preventing access comprises deferring access to the first resource at least until a set of one or more transactions recorded in a redo log for the failed node are redone, this is known in the art. An example of this is shown by Grewell, from line 25 of column 10, "Specifically, the local lock manager unit notifies the requesting instance that the specified resource is invalid. The requesting instance then performs the necessary cleanup on the resources that were held by the failed instance. Typically, such cleanup involves redoing changes that had not yet been written from volatile memory of the crashed node to the database before the crash, and removing from the database any changes made by transactions within the failed instance that had not committed at the time the instance failed. According to one embodiment, the requesting instance performs recovery on all resources that were exclusively held by the crashed instance. After the requesting

instance has performed the appropriate cleanup operations, the requesting instance sends a message to the local lock manager unit indicating that the cleanup has been completed." A person of ordinary skill in the art at the time of the invention would have been motivated to clean up after a failed master because, from Forman, "This could occur where a shadow process has commenced update but has not completed the process of sending the update to the master for propagation or where the master failed before propagating the data." And from Grewell, "The requesting instance then performs the necessary cleanup on the resources that were held by the failed instance."

Conclusion

25. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure. See notice of references cited.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Gabriel L. Chu whose telephone number is (571) 272-3656. The examiner can normally be reached on weekdays between 8:30 AM and 5:00 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Scott Baderman can be reached on (571) 272-3644. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Gabriel L. Chu
Examiner
Art Unit 2114

gc